

Disjunctive Probabilistic Modal Logic is Enough for Bisimilarity on Reactive Probabilistic Systems

Marco Bernardo

Dip. di Scienze Pure e Applicate,
Università di Urbino, Italy

Marino Miculan

Dip. di Scienze Matematiche, Informatiche e Fisiche,
Università di Udine, Italy

Abstract

Larsen and Skou characterized probabilistic bisimilarity over reactive probabilistic systems with a logic including true, negation, conjunction, and a diamond modality decorated with a probabilistic lower bound. Later on, Desharnais, Edalat, and Panangaden showed that negation is not necessary to characterize the same equivalence. In this paper, we prove that the logical characterization holds also when conjunction is replaced by disjunction, with negation still being not necessary. To this end, we introduce *reactive probabilistic trees*, a fully abstract model for reactive probabilistic systems that allows us to demonstrate expressiveness of the disjunctive probabilistic modal logic, as well as of the previously mentioned logics, by means of a compactness argument.

1. Introduction

Since its introduction by Larsen and Skou [12], *probabilistic bisimilarity* has been used to compare probabilistic systems. It corresponds to Milner's *strong bisimilarity* for nondeterministic systems, and coincides with *ordinary lumpability* for Markov chains. Larsen and Skou [12] first proved that probabilistic bisimilarity for *reactive probabilistic systems* can be given a *logical characterization*: two processes are bisimilar if and only if they satisfy the same set of formulas of a propositional modal logic similar to Hennessy-Milner logic [9]. In addition to the usual constructs \top , \neg , and \wedge , this logic features a diamond modality $\langle a \rangle_p \phi$, which is satisfied by a state if, after performing action a , the probability of being in a state satisfying ϕ is at least p .

Later on, Desharnais, Edalat, and Panangaden [5] showed that negation is *not* necessary for discrimination purposes, by working in a *continuous-state* setting. This result has no counterpart in the nonprobabilistic setting, where Hennessy-Milner logic without negation characterizes *simulation* equivalence, which is strictly coarser than bisimilarity [7] (while the two equivalences are known to coincide on reactive probabilistic processes [2]).

In this paper, we show that *disjunction* can be used in place of conjunction without having to reintroduce negation. Thus, the constructs \top , \vee , and $\langle a \rangle_p$ suffice to characterize probabilistic bisimilarity on reactive probabilistic processes. The intuition is that from a conjunctive distinguishing formula we can often derive a disjunc-

tive one by suitably increasing some probabilistic lower bounds. Not even this result has a counterpart in the nonprobabilistic setting, where replacing conjunction with disjunction in the absence of negation yields trace equivalence (this equivalence does *not* coincide with bisimilarity on reactive probabilistic processes).

The idea of the proof is the following. First, using a simple categorical construction, we show that each reactive probabilistic process can be given a semantics in a precise canonical form, which we call *reactive probabilistic tree*. These trees can be seen as the probabilistic counterpart of Winskel's synchronization trees used for nondeterministic processes. The semantics is *fully abstract*, i.e., two states are probabilistically bisimilar if and only if they are mapped to the same reactive probabilistic tree. Moreover, the semantics is *compact*, in the sense that two (possibly infinite) trees are equal if and only if all of their finite approximations are equal. Hence, in order to prove that our logic characterizes probabilistic bisimilarity, it suffices to prove that it allows to discriminate *finite* reactive probabilistic trees. Indeed, given two different finite trees, we show how to construct (by induction on the height of one of the trees) a *distinguishing formula* of the disjunctive logic that tells the two trees apart and has a depth not exceeding the height of the two trees. Our technique applies also to the logics in [5, 12], for which it allows us to provide simpler proofs of expressiveness, directly in a *discrete* setting. More generally, this technique can be used in any computational model that has a compact, fully abstract semantics.

Synopsis In Sect. 2, we recall the basic definitions about reactive probabilistic processes, bisimilarity, and logics. In Sect. 3, we characterize probabilistic bisimilarity in terms of finite reactive probabilistic trees. In Sect. 4, we prove that the various probabilistic modal logics considered in the paper can discriminate these finite trees, and hence characterize probabilistic bisimilarity. Conclusions and directions for future work are in Sect. 5.

2. Processes, Bisimilarity, and Logics

2.1 Reactive Probabilistic Processes and Strong Bisimilarity

Probabilistic processes can be represented as labeled transitions systems [11] enriched with probabilistic information used to determine which action is executed or which state is reached. Following the terminology of [8], we focus on *reactive* probabilistic processes, where every state has for each action at most one outgoing distribution over states; the choice among these arbitrarily many, differently labeled distributions is nondeterministic. For a countable (i.e., finite or countably infinite) set X , the set of finitely supported (a.k.a. simple) probability distributions over X is:

$$D(X) = \{\Delta : X \rightarrow \mathbb{R}_{[0,1]} \mid |\text{supp}(\Delta)| < \omega, \sum_{x \in X} \Delta(x) = 1\} \quad (1)$$

where the *support* is defined as $\text{supp}(\Delta) \triangleq \{x \in X \mid \Delta(x) > 0\}$.

DEFINITION 2.1. [RPLTS] A *reactive probabilistic labeled transition system*, RPLTS for short, is a triple (S, A, \longrightarrow) where:

- S is a countable set of *states*;
- A is a countable set of *actions*;
- $\longrightarrow \subseteq S \times A \times D(S)$ is a *transition relation* such that, whenever $(s, a, \Delta_1), (s, a, \Delta_2) \in \longrightarrow$, then $\Delta_1 = \Delta_2$. ■

An RPLTS can be seen as a directed graph whose edges are labeled by pairs $(a, p) \in A \times \mathbb{R}_{[0,1]}$. For every $s \in S$ and $a \in A$, if there are a -labeled edges outgoing from s , then these are finitely many (*image finiteness*), because the considered distributions are finitely supported, and the numbers on them add up to 1. As usual, we denote $(s, a, \Delta) \in \longrightarrow$ as $s \xrightarrow{a} \Delta$, where the set of reachable states coincides with $\text{supp}(\Delta)$. We also define cumulative reachability as $\Delta(S') = \sum_{s' \in S'} \Delta(s')$ for all $S' \subseteq S$.

Probabilistic bisimilarity for the class of reactive probabilistic processes was introduced by Larsen and Skou [12].

DEFINITION 2.2. [Probabilistic bisimilarity] Let (S, A, \longrightarrow) be an RPLTS. An equivalence relation \mathcal{B} over S is a *probabilistic bisimulation* iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for all actions $a \in A$:

- if $s_1 \xrightarrow{a} \Delta_1$, then $s_2 \xrightarrow{a} \Delta_2$ and $\Delta_1(C) = \Delta_2(C)$ for all equivalence classes $C \in S/\mathcal{B}$;
- if $s_2 \xrightarrow{a} \Delta_2$, then $s_1 \xrightarrow{a} \Delta_1$ and $\Delta_1(C) = \Delta_2(C)$ for all equivalence classes $C \in S/\mathcal{B}$.

We say that $s_1, s_2 \in S$ are *probabilistically bisimilar*, written $s_1 \sim_{\text{PB}} s_2$, iff there exists a probabilistic bisimulation including the pair (s_1, s_2) . ■

2.2 Probabilistic Modal Logics

In our setting, a probabilistic modal logic is a pair formed by a set \mathcal{L} of *formulas* and an RPLTS-indexed family of *satisfaction relations* $\models \subseteq S \times \mathcal{L}$. The *logical equivalence* induced by \mathcal{L} over S is defined by letting $s_1 \cong_{\mathcal{L}} s_2$, where $s_1, s_2 \in S$, iff $s_1 \models \phi \iff s_2 \models \phi$ for all $\phi \in \mathcal{L}$. We say that \mathcal{L} *characterizes* a binary relation \mathcal{R} over S when $\mathcal{R} = \cong_{\mathcal{L}}$.

We are especially interested in probabilistic modal logics characterizing \sim_{PB} . The logics considered in this paper are similar to Hennessy-Milner logic [9], but the diamond modality is decorated with a probabilistic lower bound as follows:

$$\begin{aligned} \text{PML}_{\neg\wedge} : \quad \phi &::= \top \mid \neg\phi \mid \phi \wedge \phi \mid \langle a \rangle_p \phi \\ \text{PML}_{\neg\vee} : \quad \phi &::= \top \mid \neg\phi \mid \phi \vee \phi \mid \langle a \rangle_p \phi \\ \text{PML}_{\wedge} : \quad \phi &::= \top \mid \phi \wedge \phi \mid \langle a \rangle_p \phi \\ \text{PML}_{\vee} : \quad \phi &::= \top \mid \phi \vee \phi \mid \langle a \rangle_p \phi \end{aligned}$$

where $p \in \mathbb{R}_{[0,1]}$; trailing \top 's will be omitted for sake of readability. Their semantics with respect to an RPLTS state s is as usual:

$$\begin{aligned} s \models \top &\iff \text{true} \\ s \models \neg\phi &\iff s \not\models \phi \\ s \models \phi_1 \wedge \phi_2 &\iff s \models \phi_1 \text{ and } s \models \phi_2 \\ s \models \phi_1 \vee \phi_2 &\iff s \models \phi_1 \text{ or } s \models \phi_2 \\ s \models \langle a \rangle_p \phi &\iff s \xrightarrow{a} \Delta \text{ and } \Delta(\{s' \in S \mid s' \models \phi\}) \geq p \end{aligned}$$

Larsen and Skou [12] proved that $\text{PML}_{\neg\wedge}$ characterizes \sim_{PB} . This holds true for $\text{PML}_{\neg\vee}$ as well, because $\text{PML}_{\neg\vee}$ is equivalent to $\text{PML}_{\neg\wedge}$. Desharnais, Edalat, and Panangaden [5] then proved in a *measure-theoretic* setting that PML_{\wedge} characterizes \sim_{PB} too, and hence negation is not necessary. This was subsequently redemonstrated by Jacobs and Sokolova [10] in the *dual adjunction* framework, as well as by Deng and Wu [4] with a simpler proof. The main aim of this paper is to show that PML_{\vee} suffices as well.

3. Compact Characterization of Probabilistic Bisimilarity

In this section, we provide a characterization of probabilistic bisimilarity by means of *finite* structures in a canonical form. To this end, we introduce *reactive probabilistic trees*, a concrete representation of probabilistic behaviors.

3.1 Coalgebras for Probabilistic Systems

We begin by recalling the coalgebraic setting for probabilistic systems; see, e.g., [15]. The function D defined in (1) extends to a functor $D : \text{Set} \rightarrow \text{Set}$ whose action on morphisms is, for $f : X \rightarrow Y$:

$$D(f) : D(X) \rightarrow D(Y) \quad D(f)(\Delta) = \lambda y. \Delta(f^{-1}(y))$$

Then, it is easy to see that every RPLTS corresponds to a coalgebra of the following functor:

$$B_{RP} : \text{Set} \rightarrow \text{Set} \quad B_{RP}(X) = (D(X) + 1)^A$$

Indeed, given $S = (S, A, \longrightarrow)$, we define the corresponding coalgebra (S, σ) as

$$\sigma : S \rightarrow B_{RP}(S) \quad \sigma(s) \triangleq \lambda a. \begin{cases} \Delta & \text{if } s \xrightarrow{a} \Delta \\ * & \text{otherwise} \end{cases}$$

A *homomorphism* $h : (S, \sigma) \rightarrow (T, \tau)$ is a function $h : S \rightarrow T$ which respects the coalgebraic structures, i.e., $\tau \circ h = (B_{RP}h) \circ \sigma$. We denote by $\text{Coalg}(B_{RP})$ the category of B_{RP} -coalgebras and their homomorphisms.

Azel and Mendler [1] introduced a general notion of bisimulation for coalgebras, which in our setting instantiates as follows:

DEFINITION 3.1. Let (S_1, σ_1) and (S_2, σ_2) be B_{RP} -coalgebras. A relation $\mathcal{R} \subseteq S_1 \times S_2$ is a *B_{RP} -bisimulation* iff there exists a coalgebra structure $\rho : \mathcal{R} \rightarrow B_{RP}\mathcal{R}$ such that the projections $\pi_1 : \mathcal{R} \rightarrow S_1$ and $\pi_2 : \mathcal{R} \rightarrow S_2$ are homomorphisms (i.e., $\sigma_i \circ \pi_i = B_{RP}\pi_i \circ \rho$ for $i = 1, 2$).

We say that $s_1 \in S_1$ and $s_2 \in S_2$ are *B_{RP} -bisimilar*, written $s_1 \sim s_2$, iff there exists a B_{RP} -bisimulation including (s_1, s_2) . ■

The following result shows that probabilistic bisimilarity corresponds to B_{RP} -bisimilarity.

PROPOSITION 3.2. The probabilistic bisimilarity over an RPLTS (S, A, \longrightarrow) coincides with the B_{RP} -bisimilarity over the corresponding coalgebra (S, σ) .

PROOF An immediate consequence of [15, Lemma 4.4 and Thm. 4.5]. ■

The next step is to associate each state of a given RPLTS with its *behavior*, i.e., a structure in some canonical form which we can reason about. These structures can be seen as the elements of the final coalgebra of B_{RP} , which exists because we consider only finitely supported distributions, as proved in [15, Thm. 4.6]:

PROPOSITION 3.3. The functor B_{RP} admits final coalgebra.

PROOF The functor D is bounded because it is restricted to distributions with finite support. Hence also B_{RP} is bounded; then the final coalgebra exists by the general result [14, Thm. 10.4]. ■

Let (Z, ζ) be a final B_{RP} -coalgebra (which is unique up-to isomorphism). This coalgebra can be seen as the RPLTS which subsumes all possible behaviors of any RPLTS. Moreover, elements of Z can be seen as “canonical” representatives of behaviors, because different states of Z are never bisimilar:

PROPOSITION 3.4. For all $z_1, z_2 \in Z$: $z_1 \sim z_2$ iff $z_1 = z_2$. ■

3.2 Reactive Probabilistic Trees

Although Prop. 3.3 guarantees the existence of the final coalgebra, it does not provide us with a concrete representation of its elements. In this subsection, we introduce *reactive probabilistic trees*, a representation of the final B_{RP} -coalgebra which can be seen as the natural extension to the probabilistic setting of *strongly extensional trees* used to represent the final \mathcal{P}_f -coalgebra [16].

DEFINITION 3.5. [RPT] An (A -labeled) *reactive probabilistic tree* is a pair $(X, succ)$ where $X \in \mathbf{Set}$ and $succ : X \times A \rightarrow \mathcal{P}_f(X \times \mathbb{R}_{(0,1]})$ are such that the relation \leq over X defined by:

$$\frac{x \leq y \quad z \in succ(y, a)}{x \leq z}$$

is a partial order with a least element, called *root*, and for all $x \in X$ and $a \in A$:

1. the set $\{y \in X \mid y \leq x\}$ is finite and well-ordered;
2. for all $(x_1, p_1), (x_2, p_2) \in succ(x, a)$, if $x_1 = x_2$ then $p_1 = p_2$;
3. for all $(x_1, p_1), (x_2, p_2) \in succ(x, a)$, if the subtrees rooted at x_1 and x_2 are isomorphic then $x_1 = x_2$;
4. if $succ(x, a) \neq \emptyset$ then $\sum_{(y,p) \in succ(x,a)} p = 1$.

We denote by RPT , ranged over by t, t_1, t_2, \dots , the set of all reactive probabilistic trees (possibly of infinite height), up-to isomorphism. ■

Reactive probabilistic trees are unordered trees where each node for each action has either no successor or a finite set of successors labeled with a positive real number adding up to 1; moreover, subtrees rooted at these successors are all different. See the forthcoming Fig. 1 for some examples. In particular, the trivial tree is $nil \triangleq (\{\perp\}, \lambda x, a. \emptyset)$.

For $t = (X, succ)$, we denote its root by \perp_t , its a -successors by $t(a) \triangleq succ(\perp_t, a)$, and the subtree rooted at $x \in X$ by $t[x] \triangleq (\{y \in X \mid x \leq y\}, \lambda y, a. succ(y, a))$; thus, $\perp_{t[x]} = x$.

We define $height : RPT \rightarrow \mathbb{N} \cup \{\omega\}$ in the obvious way:

$$height(t) \triangleq \sup\{1 + height(t') \mid (t', p) \in t(a), a \in A\}$$

where $\sup \emptyset = 0$; hence, $height(nil) = 0$. In particular, we denote by $RPT_f \triangleq \{t \in RPT \mid height(t) < \omega\}$ the set of reactive probabilistic trees of finite height.

A (possibly infinite) tree can be truncated at any height n , yielding a finite tree where the missing subtrees are replaced by nil . In order to obtain a RPT_f , we need to collapse isomorphic subtrees resulting from the truncation. More formally, we define first the *truncation* function tr_n by induction on n : $tr_0(t) \triangleq nil$ and

$$tr_{n+1}(t) \triangleq (\{\perp_t\} \cup \bigcup \{X' \mid ((X', succ'), p') \in tr_n(t(a)), a \in A\}, succ_Y)$$

where $succ_Y(\perp_t, a) \triangleq \{(\perp_{t'}, p') \mid (t', p') \in q(t(a))\}$.

The tree returned by tr_n is always finite, but possibly not extensional. Hence we have to collapse its isomorphic subtrees adding up their weights by means of the *coll* function as follows:

$$coll(t) \triangleq (\{\perp_t\} \cup \bigcup \{X' \mid ((X', succ'), p') \in U_a, a \in A\}, succ_c)$$

where $W_a = \{coll(t'), p \mid (t', p) \in succ(\perp_t, a)\}$

$$U_a = \{(s, \sum_{(s,p) \in W_a} p) \mid s \in \pi_1(W_a)\}$$

$$succ_c = (\bigcup_{a \in A} \{(\perp_t, a) \mapsto \{(s, p) \mid (s, p) \in U_a\}\}) \cup \bigcup_{(s,p) \in U_a} succ_s$$

Finally we can define the *pruning* of t as $t|_n \triangleq coll(tr_n(t))$.

We have now to show that RPT is (the carrier of) the final B_{RP} -coalgebra (up-to isomorphism). In order to simplify the proof, we reformulate B_{RP} in a slightly more “relational” format. We define a functor $D' : \mathbf{Set} \rightarrow \mathbf{Set}$ by letting for any set X :

$$D'X = \{U \in \mathcal{P}_f(X \times \mathbb{R}_{(0,1]}) \mid \text{if } U \neq \emptyset \text{ then } \sum_{(x,p) \in U} p = 1 \text{ and } \forall (x, p), (x', p') \in U : x = x' \Rightarrow p = p'\}$$

and for any $f : X \rightarrow Y$, the function $D'f : D'X \rightarrow D'Y$ maps $U \in D'X$ to $\{(f(x), \sum_{(x,p) \in U} p) \mid x \in \pi_1(U)\}$. Then:

PROPOSITION 3.6. 1. $D' \cong D + 1$.

2. $D'^A \cong B_{RP}$.

3. $Coalg(D'^A) \cong Coalg(B_{RP})$.

4. The supports of the final D'^A -coalgebra and of the final B_{RP} -coalgebra are isomorphic.

PROOF 1: For $X \in \mathbf{Set}$, define $\phi_X : D'X \rightarrow DX + 1$ as $\phi_X(\emptyset) = *$, and for $U \neq \emptyset$, $\phi_X(U) : X \rightarrow \mathbb{R}_{[0,1]}$ maps x to p if $(x, p) \in U$, to 0 otherwise. It is easy to check that the ϕ_X 's are invertible and form a natural isomorphism $\phi : D' \xrightarrow{\sim} D + 1$.

2: Trivial by 1; let $\psi : D'^A \xrightarrow{\sim} B_{RP}$ be the underlying natural isomorphism.

3: A D'^A -coalgebra $(X, \sigma : X \rightarrow D'(X)^A)$ is mapped to $(X, \psi_X \circ \sigma : X \rightarrow B_{RP}(X))$; the vice versa is similar, using ψ_X^{-1} . It is easy to check that these maps are inverse to each other.

4: Trivial by 3. ■

We can now prove that RPT is the carrier of the final B_{RP} -coalgebra (up-to isomorphism). First, we observe that the set RPT can be endowed with a D'^A -coalgebra structure $\rho : RPT \rightarrow (D'(RPT))^A$ defined as follows, for $t = (X, succ)$:

$$\rho(t)(a) \triangleq \{(t[x], p) \mid (x, p) \in succ(\perp_t, a)\}$$

THEOREM 3.7. (RPT, ρ) is a final B_{RP} -coalgebra.

PROOF By Prop. 3.6, it suffices to prove that (RPT, ρ) is the final D'^A -coalgebra. To this end, we follow the construction given by Worrell in [16, Thm. 11]. We define an ordinal-indexed *final sequence* of sets $(B_\alpha)_\alpha$ together with “projection functions” $(f_\gamma^\beta : B_\beta \rightarrow B_\gamma)_{\gamma \leq \beta}$:

$$\begin{aligned} B_0 &= \{nil\} \cong 1 & f_0^1 &= ! \\ B_{\alpha+1} &= D'(B_\alpha)^A & f_{\alpha+1}^{\alpha+2} &= D'(f_\alpha^{\alpha+1})^A \\ B_\lambda &= \lim_{\alpha < \lambda} B_\alpha & f_\alpha^\lambda &= \pi_\alpha \quad \text{for } \lambda \text{ a limit ordinal} \end{aligned}$$

the remaining f_γ^β being given by suitable compositions. D is ω -accessible (because we restrict to finitely supported distributions), thus by [16, Thm. 13] and Prop. 3.6 the final sequence converges in at most $\omega + \omega$ steps to the set $B_{\omega+\omega}$ which is the carrier of the final D'^A -coalgebra.

Now, we have to prove that $B_{\omega+\omega}$ is isomorphic to RPT . An element of $B_{\omega+\omega}$ is a sequence of finite trees $\vec{t} = (t_0, t_1, \dots)$ such that for each $k \in \omega$ there exists $N_k \in \mathbb{N}$ such that nodes at depth k of any tree t_i have at most N_k successors for each label $a \in A$. These sequences can be seen as compatible partial views of a single (possibly infinite) tree. Thus, given a sequence \vec{t} the corresponding tree $u \in RPT$ is obtained by *amalgamating* \vec{t} : u at depth k is defined by the level k of a suitable tree t_i , where i is such that for all $j \geq i$, t_j is equal to t_i up to depth k . On the other hand, given $u \in RPT$ we can define the corresponding sequence $\vec{t} \in B_{\omega+\omega}$ as $t_i = u|_i$.

It can be checked that these two maps form an isomorphism between $B_{\omega+\omega}$ and RPT . Moreover, they respect the coalgebraic structures, where $\tau : B_{\omega+\omega} \rightarrow D'(B_{\omega+\omega})^A$ is given by $\tau(\vec{t})(a) = \{\vec{t}' \in B_{\omega+\omega} \mid \forall i \in \omega : t'_i \in succ(t_i, a)\}$. Therefore, $(B_{\omega+\omega}, \tau)$ and (RPT, ρ) are isomorphic D'^A -coalgebras, hence the thesis. ■

3.3 Full Abstraction and Compactness

By virtue of Thm. 3.7, given an RPLTS $S = (S, A, \longrightarrow)$ there exists a unique coalgebra homomorphism $\llbracket \cdot \rrbracket : S \rightarrow RPT$, called the (final) semantics of S , which associates each state in S with its behavior. This semantics is *fully abstract*.

THEOREM 3.8. [Full abstraction] Let (S, A, \longrightarrow) be an RPLTS. For all $s_1, s_2 \in S$: $s_1 \sim_{PB} s_2$ iff $\llbracket s_1 \rrbracket = \llbracket s_2 \rrbracket$.

PROOF It follows from Props. 3.2 and 3.4 and Thm. 3.7. ■

A key property of reactive probabilistic trees is that they are *compact*: two different trees can be distinguished by looking at their finite subtrees only. Let us formalize this principle:

THEOREM 3.9. [Compactness] For all $t_1, t_2 \in RPT$: $t_1 = t_2$ iff for all $n \in \mathbb{N}$: $t_1|_n = t_2|_n$.

PROOF The “only if” is trivial. For the “if” direction, let us assume that $t_1 \neq t_2$; we have to find n such that $t_1|_n \neq t_2|_n$. Given a tree u_0 , a *finite path* in u_0 is a sequence $(a_1, p_1, a_2, p_2, \dots, a_n, p_n)$ such that for $i = 1, \dots, n$: $(x_i, p_i) \in u_{i-1}(a_i)$ and $u_i = u[x_i]$. If $t_1 \neq t_2$, then there is a path of length n in, say, t_1 which cannot be replayed in t_2 : in t_2 we reach a tree t'_{n-1} such that for all t , $(t, p_n) \notin t'_{n-1}(a_n)$. Therefore $t_1|_n \neq t_2|_n$. ■

COROLLARY 3.10. Let (S, A, \longrightarrow) be an RPLTS. For all $s_1, s_2 \in S$: $s_1 \sim_{PB} s_2$ iff for all $n \in \mathbb{N}$: $\llbracket s_1 \rrbracket|_n = \llbracket s_2 \rrbracket|_n$. ■

4. The Discriminating Power of PML_{\vee}

By virtue of the categorical construction leading to Cor. 3.10, in order to prove that a modal logic characterizes \sim_{PB} over reactive probabilistic processes, it is enough to show that it can discriminate all reactive probabilistic trees of *finite* height. A specific condition on the depth of distinguishing formulas has also to be satisfied, where $depth(\phi)$ is defined as usual:

$$\begin{aligned} depth(\top) &= 0 \\ depth(\neg\phi') &= depth(\phi') \\ depth(\phi_1 \wedge \phi_2) &= \max(depth(\phi_1), depth(\phi_2)) \\ depth(\phi_1 \vee \phi_2) &= \max(depth(\phi_1), depth(\phi_2)) \\ depth(\langle a \rangle_p \phi') &= 1 + depth(\phi') \end{aligned}$$

PROPOSITION 4.1. Let \mathcal{L} be one of the probabilistic modal logics in Sect. 2.2. If \mathcal{L} characterizes $=$ over RPT_f and for any two nodes t_1 and t_2 of an arbitrary RPT_f model such that $t_1 \neq t_2$ there exists $\phi \in \mathcal{L}$ distinguishing t_1 from t_2 such that:

$$depth(\phi) \leq \max(height(t_1), height(t_2))$$

then \mathcal{L} characterizes \sim_{PB} over the set of RPLTS models.

PROOF Given two states s_1 and s_2 of an RPLTS, if $s_1 \sim_{PB} s_2$ then for all $n \in \mathbb{N}$ it holds that $\llbracket s_1 \rrbracket|_n = \llbracket s_2 \rrbracket|_n$ thanks to Cor. 3.10, hence s_1 and s_2 satisfy the same formulas of \mathcal{L} because \mathcal{L} characterizes $=$ over RPT_f . Suppose now that $s_1 \not\sim_{PB} s_2$ and consider the minimum $n \in \mathbb{N}_{\geq 1}$ for which $\llbracket s_1 \rrbracket|_n \neq \llbracket s_2 \rrbracket|_n$. Then there exists $\phi \in \mathcal{L}$ distinguishing $\llbracket s_1 \rrbracket|_n$ from $\llbracket s_2 \rrbracket|_n$ such that $depth(\phi) \leq \max(height(\llbracket s_1 \rrbracket|_n), height(\llbracket s_2 \rrbracket|_n)) = n$, hence the same formula ϕ also distinguishes s_1 from s_2 . ■

Notice that, in the proof above, if $depth(\phi)$ were greater than n then, in general, ϕ may not distinguish higher prunings of $\llbracket s_1 \rrbracket$ and $\llbracket s_2 \rrbracket$, nor may any formula of depth at most n and derivable from ϕ still distinguish $\llbracket s_1 \rrbracket|_n$ from $\llbracket s_2 \rrbracket|_n$.

EXAMPLE 4.2. Consider a process whose initial state s_1 has only an a -transition to a state having only a c -transition to nil , and another process whose initial state s_2 has only a b -transition to a

state having only a d -transition to nil . Their corresponding trees differ at height $n = 1$ because $\llbracket s_1 \rrbracket|_1$ has an a -transition to nil while $\llbracket s_2 \rrbracket|_1$ has a b -transition to nil .

The formula of depth 2 given by $\langle a \rangle_1 \neg \langle c \rangle_1$ distinguishes $\llbracket s_1 \rrbracket|_1$ from $\llbracket s_2 \rrbracket|_1$, but this is no longer the case with $\llbracket s_1 \rrbracket|_2$ and $\llbracket s_2 \rrbracket|_2$ as neither satisfies that formula.

The formula of depth 2 given by $\langle a \rangle_1 \vee \langle b \rangle_1 \langle c \rangle_1$ distinguishes $\llbracket s_1 \rrbracket|_1$ from $\llbracket s_2 \rrbracket|_1$, but this is no longer the case with the derived formula $\langle a \rangle_1 \vee \langle b \rangle_1$ of depth 1 as both nodes satisfy it. ■

Based on the considerations above, in this section we show the main result of the paper: the logical equivalence induced by PML_{\vee} has the same discriminating power as \sim_{PB} .

This result is accomplished in three steps. Firstly, we redemonstrate Larsen and Skou’s result for PML_{\wedge} in the RPT_f setting, which yields a proof that, with respect to the one in [12], is simpler and does not require the minimal deviation assumption (i.e., that the probability associated with any state in the support of the target distribution of a transition be a multiple of some value). This provides a proof scheme for the subsequent steps. Secondly, we demonstrate that PML_{\vee} characterizes \sim_{PB} by adapting the proof scheme to cope with the replacement of conjunction with disjunction. Thirdly, we demonstrate that PML_{\vee} characterizes \sim_{PB} by further adapting the proof scheme to cope with the absence of negation.

Moreover, we redemonstrate Desharnais, Edalat, and Panangaden’s result for PML_{\wedge} through yet another adaptation of the proof scheme that, unlike the proof in [5], is constructive and works directly on *discrete* state spaces without making use of measure-theoretic arguments based on analytic spaces. Avoiding the resort to measure theory was shown to be possible for the first time by Worrell in an unpublished note cited in [13].

4.1 PML_{\wedge} Characterizes \sim_{PB} : A New Proof

To show that the logical equivalence induced by PML_{\wedge} implies node equality $=$, we reason on the contrapositive. Given two nodes t_1 and t_2 such that $t_1 \neq t_2$, we proceed by induction on the height of t_1 to find a distinguishing PML_{\wedge} formula whose depth is not greater than the heights of t_1 and t_2 . The idea is to exploit negation, so to ensure that certain distinguishing formulas are *satisfied* by a certain derivative t' of t_1 (rather than the derivatives of t_2 different from t'), then take the *conjunction* of those formulas preceded by a diamond decorated with the probability for t_1 of *reaching* t' .

The only non-trivial case is the one in which t_1 and t_2 enable the same actions. At least one of those actions, say a , is such that, after performing it, the two nodes reach two distributions $\Delta_{1,a}$ and $\Delta_{2,a}$ such that $\Delta_{1,a} \neq \Delta_{2,a}$. Given a node $t' \in \text{supp}(\Delta_{1,a})$ such that $\Delta_{1,a}(t') > \Delta_{2,a}(t')$, by the induction hypothesis there exists a PML_{\wedge} formula $\phi'_{2,j}$ that distinguishes t' from a specific $t'_{2,j} \in \text{supp}(\Delta_{2,a}) \setminus \{t'\}$. We can assume that $t' \models \phi'_{2,j} \not\models t'_{2,j}$ otherwise, thanks to the presence of negation in PML_{\wedge} , it would suffice to consider $\neg\phi'_{2,j}$.

As a consequence, $t_1 \models \langle a \rangle_{\Delta_{1,a}(t')} \bigwedge_j \phi'_{2,j} \not\models t_2$ because $\Delta_{1,a}(t') > \Delta_{2,a}(t')$ and $\Delta_{2,a}(t')$ is the maximum probabilistic lower bound for which t_2 satisfies a formula of that form. Notice that $\Delta_{1,a}(t')$ may not be the maximum probabilistic lower bound for which t_1 satisfies such a formula, because $\bigwedge_j \phi'_{2,j}$ might be satisfied by other a -derivatives of t_1 in $\text{supp}(\Delta_{1,a}) \setminus \{t'\}$.

THEOREM 4.3. Let (T, A, \longrightarrow) be in RPT_f and $t_1, t_2 \in T$. Then $t_1 = t_2$ iff $t_1 \models \phi \iff t_2 \models \phi$ for all $\phi \in PML_{\wedge}$. Moreover, if $t_1 \neq t_2$, then there exists $\phi \in PML_{\wedge}$ distinguishing t_1 from t_2 such that $depth(\phi) \leq \max(height(t_1), height(t_2))$.

PROOF Given $t_1, t_2 \in T$, we proceed as follows:

- If $t_1 = t_2$, then obviously $t_1 \models \phi \iff t_2 \models \phi$ for all $\phi \in PML_{\wedge}$.

- Assuming that $t_1 \neq t_2$, we show that there exists $\phi \in \text{PML}_{\neg\wedge}$, with $\text{depth}(\phi) \leq \max(\text{height}(t_1), \text{height}(t_2))$, such that it is not the case that $t_1 \models \phi \iff t_2 \models \phi$ by proceeding by induction on $\text{height}(t_1) \in \mathbb{N}$:
 - If $\text{height}(t_1) = 0$, then $\text{height}(t_2) \geq 1$ because $t_1 \neq t_2$. As a consequence, t_2 has at least one outgoing transition, say labeled with a , hence $t_1 \not\models \langle a \rangle_1 \models t_2$. Notice that $\text{depth}(\langle a \rangle_1) = 1 \leq \max(\text{height}(t_1), \text{height}(t_2))$.
 - Let $\text{height}(t_1) = n + 1$ for some $n \in \mathbb{N}$ and suppose that for all $t'_1, t'_2 \in T$ such that $t'_1 \neq t'_2$ and $\text{height}(t'_1) \leq n$ there exists $\phi' \in \text{PML}_{\neg\wedge}$, with $\text{depth}(\phi') \leq \max(\text{height}(t'_1), \text{height}(t'_2))$, such that it is not the case that $t'_1 \models \phi' \iff t'_2 \models \phi'$. Let $\text{init}(t_h)$, $h \in \{1, 2\}$, be the set of actions in A labeling the transitions departing from t_h :
 - If $\text{init}(t_1) \neq \text{init}(t_2)$, then it holds that $t_1 \models \langle a \rangle_1 \not\models t_2$ for some $a \in \text{init}(t_1) \setminus \text{init}(t_2)$ or $t_1 \not\models \langle a \rangle_1 \models t_2$ for some $a \in \text{init}(t_2) \setminus \text{init}(t_1)$. Notice that $\text{depth}(\langle a \rangle_1) = 1 \leq \max(\text{height}(t_1), \text{height}(t_2))$.
 - If $\text{init}(t_1) = \text{init}(t_2)$, then $\text{init}(t_1) \neq \emptyset \neq \text{init}(t_2)$ as $\text{height}(t_1) \geq 1$. Since $t_1 \neq t_2$, there must exist $a \in \text{init}(t_1)$ such that $t_1 \xrightarrow{a} \Delta_{1,a}$, $t_2 \xrightarrow{a} \Delta_{2,a}$, and $\Delta_{1,a} \neq \Delta_{2,a}$. From $\Delta_{1,a} \neq \Delta_{2,a}$, it follows that there exists $t' \in \text{supp}(\Delta_{1,a})$ such that $1 \geq \Delta_{1,a}(t') > \Delta_{2,a}(t') \geq 0$. Assuming that $\text{supp}(\Delta_{2,a}) \setminus \{t'\} = \{t'_{2,1}, t'_{2,2}, \dots, t'_{2,k}\}$, which cannot be empty because there must also exist $t'_2 \in \text{supp}(\Delta_{2,a})$ such that $0 \leq \Delta_{1,a}(t'_2) < \Delta_{2,a}(t'_2) \leq 1$, by the induction hypothesis for each $j = 1, 2, \dots, k$ there exists $\phi'_{2,j} \in \text{PML}_{\neg\wedge}$, with $\text{depth}(\phi'_{2,j}) \leq \max(\text{height}(t'), \text{height}(t'_{2,j}))$, such that it is not the case that $t' \models \phi'_{2,j} \iff t'_{2,j} \models \phi'_{2,j}$. Since $\text{PML}_{\neg\wedge}$ includes negation, without loss of generality we can assume that $t' \models \phi'_{2,j} \not\models t'_{2,j}$. Therefore, it holds that $t_1 \models \langle a \rangle_{\Delta_{1,a}(t')} \bigwedge_{1 \leq j \leq k} \phi'_{2,j} \not\models t_2$ because $\Delta_{1,a}(t') > \Delta_{2,a}(t')$ and $\Delta_{2,a}(t')$ is the maximum probabilistic lower bound for which t_2 satisfies a formula of that form. Notice that the resulting formula, which we denote by ϕ for short, satisfies:

$$\begin{aligned} \text{depth}(\phi) &= 1 + \max_{1 \leq j \leq k} \text{depth}(\phi'_{2,j}) \\ &\leq 1 + \max_{1 \leq j \leq k} \max(\text{height}(t'), \text{height}(t'_{2,j})) \\ &= 1 + \max(\text{height}(t'), \max_{1 \leq j \leq k} \text{height}(t'_{2,j})) \\ &= \max(1 + \text{height}(t'), 1 + \max_{1 \leq j \leq k} \text{height}(t'_{2,j})) \\ &\leq \max(\text{height}(t_1), \text{height}(t_2)) \end{aligned}$$

4.2 $\text{PML}_{\neg\vee}$ Characterizes \sim_{PB} : Adapting the Proof

Since $\phi_1 \wedge \phi_2$ is logically equivalent to $\neg(\neg\phi_1 \vee \neg\phi_2)$, it is not surprising that $\text{PML}_{\neg\vee}$ characterizes \sim_{PB} too. However, the proof of this result will be useful to set up an outline of the proof of the main result of this paper, i.e., that PML_{\vee} characterizes \sim_{PB} as well.

Similarly to Thm. 4.3, also for $\text{PML}_{\neg\vee}$ we reason on the contrapositive and proceed by induction. Given t_1, t_2 such that $t_1 \neq t_2$, we intend to exploit negation, so to ensure that certain distinguishing formulas are *not satisfied* by a certain derivative t' of t_1 (rather than the derivatives of t_2 different from t'), then take the *disjunction* of those formulas preceded by a diamond decorated with the probability for t_2 of *not reaching* t' . The only non-trivial case is $t' \in \text{supp}(\Delta_{1,a})$ such that $\Delta_{1,a}(t') > \Delta_{2,a}(t')$. By inductive hypothesis there exists a formula $\phi'_{2,j}$ distinguishing t' from a specific $t'_{2,j} \in \text{supp}(\Delta_{2,a}) \setminus \{t'\}$. We can assume that $t' \not\models \phi'_{2,j} \models t'_{2,j}$ (otherwise we consider $\neg\phi'_{2,j}$ since negation is in $\text{PML}_{\neg\vee}$). Therefore, $t_1 \not\models \langle a \rangle_{1-\Delta_{2,a}(t')} \bigvee_j \phi'_{2,j} \models t_2$ because $1 - \Delta_{2,a}(t') > 1 - \Delta_{1,a}(t')$ and the maximum probabilistic lower bound for which t_1 satisfies a formula of that form cannot exceed $1 - \Delta_{1,a}(t')$. Notice that $1 - \Delta_{2,a}(t')$ is the *maximum* lower bound for which t_2 satisfies such a formula, because that value is the probability with which t_2 does not reach t' after performing a .

THEOREM 4.4. Let (T, A, \longrightarrow) be in RPT_f and $t_1, t_2 \in T$. Then $t_1 = t_2$ iff $t_1 \models \phi \iff t_2 \models \phi$ for all $\phi \in \text{PML}_{\neg\vee}$. Moreover, if $t_1 \neq t_2$, then there exists $\phi \in \text{PML}_{\neg\vee}$ distinguishing t_1 from t_2 such that $\text{depth}(\phi) \leq \max(\text{height}(t_1), \text{height}(t_2))$.

PROOF The proof is similar to the one of Thm. 4.3, apart from the final part of the last subcase, which changes as follows.

By the induction hypothesis, for each $j = 1, 2, \dots, k$ there exists $\phi'_{2,j} \in \text{PML}_{\neg\vee}$, with $\text{depth}(\phi'_{2,j}) \leq \max(\text{height}(t'), \text{height}(t'_{2,j}))$, such that it is not the case that $t' \models \phi'_{2,j} \iff t'_{2,j} \models \phi'_{2,j}$. Since $\text{PML}_{\neg\vee}$ includes negation, without loss of generality we can assume that $t' \not\models \phi'_{2,j} \models t'_{2,j}$. Therefore, it holds that $t_1 \not\models \langle a \rangle_{1-\Delta_{2,a}(t')} \bigvee_{1 \leq j \leq k} \phi'_{2,j} \models t_2$ because $1 - \Delta_{2,a}(t') > 1 - \Delta_{1,a}(t')$ and the maximum probabilistic lower bound for which t_1 satisfies a formula of that form cannot exceed $1 - \Delta_{1,a}(t')$. ■

4.3 Also PML_{\vee} Characterizes \sim_{PB}

The proof that PML_{\vee} characterizes \sim_{PB} is inspired by the one for $\text{PML}_{\neg\vee}$, thus considers the contrapositive and proceeds by induction. In the only non-trivial case, we will arrive at a situation in which $t_1 \not\models \langle a \rangle_{1-(\Delta_{2,a}(t') + p)} \bigvee_{j \in J} \phi'_{2,j} \models t_2$ for:

- a derivative t' of t_1 , such that $\Delta_{1,a}(t') > \Delta_{2,a}(t')$, not satisfying any subformula $\phi'_{2,j}$;
- a suitable probabilistic value p such that $\Delta_{2,a}(t') + p < 1$;
- an index set J identifying certain derivatives of t_2 other than t' .

The choice of t' is crucial, because negation is no longer available in PML_{\vee} . Unlike the case of $\text{PML}_{\neg\vee}$, this induces the introduction of p and the limitation to J in the format of the distinguishing formula. An important observation is that, in many cases, a disjunctive distinguishing formula can be obtained from a conjunctive one by suitably *increasing* some probabilistic lower bounds.

EXAMPLE 4.5. The nodes t_1 and t_2 in Fig. 1(a) cannot be distinguished by any formula in which neither conjunction nor disjunction occurs. It holds that:

$$\begin{aligned} t_1 &\models \langle a \rangle_{0.5} (\langle b \rangle_1 \wedge \langle c \rangle_1) \not\models t_2 \\ t_1 &\not\models \langle a \rangle_{1.0} (\langle b \rangle_1 \vee \langle c \rangle_1) \models t_2 \end{aligned}$$

Notice that, when moving from the conjunctive formula to the disjunctive one, the probabilistic lower bound decorating the a -diamond increases from 0.5 to 1 and the roles of t_1 and t_2 with respect to \models are inverted.

The situation is similar for the nodes t_3 and t_4 in Fig. 1(b), where two occurrences of conjunction/disjunction are necessary:

$$\begin{aligned} t_3 &\models \langle a \rangle_{0.2} (\langle b \rangle_1 \wedge \langle c \rangle_1 \wedge \langle d \rangle_1) \not\models t_4 \\ t_3 &\models \langle a \rangle_{0.9} (\langle b \rangle_1 \vee \langle c \rangle_1 \vee \langle d \rangle_1) \not\models t_4 \end{aligned}$$

but the roles of t_3 and t_4 with respect to \models cannot be inverted. ■

However, increasing some of the probabilistic lower bounds in a conjunctive distinguishing formula does not always yield a disjunctive one. This is the case when the use of conjunction/disjunction is not necessary for telling two different nodes apart.

EXAMPLE 4.6. For the nodes t_5 and t_6 in Fig. 1(c), it holds that:

$$t_5 \not\models \langle a \rangle_{0.5} (\langle b \rangle_1 \wedge \langle c \rangle_1) \models t_6$$

If we replace conjunction with disjunction and we vary the probabilistic lower bound between 0.5 and 1, we produce no disjunctive formula capable of discriminating between t_5 and t_6 . Nevertheless, a distinguishing formula belonging to PML_{\vee} exists because:

$$t_5 \not\models \langle a \rangle_{0.5} \langle b \rangle_1 \models t_6$$

where disjunction does not occur at all. ■




Figure 1. RPT_f models used in the examples of Sects. 4.3 and 4.4

The examples above show that the increase of some probabilistic lower bounds when moving from conjunctive distinguishing formulas to disjunctive ones takes place only in the case that the probabilities of reaching certain nodes have to be *summed up*. Additionally, we recall that, in order for two nodes to be related by \sim_{PB} , they must enable the same actions, so focussing on a *single* action is enough for discriminating when only disjunction is available.

Bearing this in mind, for any node t of finite height we define the set $\Phi_{\vee}(t)$ of PML_{\vee} formulas satisfied by t featuring:

- probabilistic lower bounds of diamonds that are *maximal* with respect to the satisfiability of a formula of that format by t (this is consistent with the observation in the last sentence before Thm. 4.4, and keeps the set $\Phi_{\vee}(t)$ finite);

- diamonds that arise only from *existing* transitions that depart from t (so to avoid useless diamonds in disjunctions and hence keep the set $\Phi_{\vee}(t)$ finite);
- disjunctions that stem only from *single* transitions of *different* nodes in the support of a distribution reached by t (transitions departing from the same node would result in formulas like $\bigvee_{h \in H} \langle a_h \rangle_{p_h} \phi_h$, with $a_{h_1} \neq a_{h_2}$ for $h_1 \neq h_2$, which are useless for discriminating with respect to \sim_{PB}) and are preceded by a diamond decorated with the *sum* of the probabilities assigned to those nodes by the distribution reached by t .

DEFINITION 4.7. The set $\Phi_{\vee}(t)$ for a node t of finite height is defined by induction on $height(t)$ as follows:

- If $height(t) = 0$, then $\Phi_{\vee}(t) = \emptyset$.

- If $\text{height}(t) \geq 1$ for t having transitions of the form $t \xrightarrow{a_i} \Delta_i$ with $\text{supp}(\Delta_i) = \{t'_{i,j} \mid j \in J_i\}$ and $i \in I \neq \emptyset$, then:

$$\begin{aligned} \Phi_{\vee}(t) &= \{\langle a_i \rangle_1 \mid i \in I\} \\ &\cup \bigcup_{i \in I} \text{hplb} \left(\bigcup_{\emptyset \neq J' \subseteq J_i} \{ \langle a_i \rangle_{\sum_{j \in J'} \Delta_i(t'_{i,j})} \bigvee_{j \in J'} \phi'_{i,j,k} \mid \right. \\ &\quad \left. t'_{i,j} \in \text{supp}(\Delta_i), \phi'_{i,j,k} \in \Phi_{\vee}(t'_{i,j}) \} \right) \end{aligned}$$

where operator \bigvee is a variant of \vee in which identical operands are not admitted (i.e., idempotence is forced) and function *hplb* keeps only the formula with the highest probabilistic lower bound decorating the initial a_i -diamond among the formulas differing only for that bound. ■

To illustrate the definition given above, we exhibit some examples showing the usefulness of Φ_{\vee} -sets for discrimination purposes. In particular, let us reconsider the non-trivial case mentioned at the beginning of this subsection. Given two different nodes that with the same action reach two different distributions, a good criterion for choosing t' (a derivative of the first node not satisfying certain formulas, to which the first distribution assigns a probability greater than the second one) seems to be the *minimality* of the Φ_{\vee} -set.

EXAMPLE 4.8. For the nodes t_7 and t_8 in Fig. 1(d), we have:

$$\begin{aligned} \Phi_{\vee}(t_7) &= \{\langle a \rangle_1, \langle a \rangle_1 \langle b \rangle_1\} \\ \Phi_{\vee}(t_8) &= \{\langle a \rangle_1, \langle a \rangle_1 \langle b \rangle_1, \langle a \rangle_1 \langle c \rangle_1\} \end{aligned}$$

A formula like $\langle a \rangle_1 (\langle b \rangle_1 \vee \langle c \rangle_1)$ is useless for discriminating between t_7 and t_8 , because disjunction is between two actions enabled by the same node and hence constituting a nondeterministic choice. Indeed, such a formula is not part of $\Phi_{\vee}(t_8)$. While in the case of conjunction it is often necessary to concentrate on several alternative actions, in the case of disjunction it is convenient to focus on a single action per node when aiming at producing a distinguishing formula.

The fact that $\langle a \rangle_1 \langle c \rangle_1 \in \Phi_{\vee}(t_8)$ is a distinguishing formula can be retrieved as follows. Starting from the two identically labeled transitions $t_7 \xrightarrow{a} \Delta_{7,a}$ and $t_8 \xrightarrow{a} \Delta_{8,a}$ where $\Delta_{7,a}(t'_7) = 1 = \Delta_{8,a}(t'_8)$ and $\Delta_{7,a}(t'_8) = 0 = \Delta_{8,a}(t'_7)$, we have:

$$\begin{aligned} \Phi_{\vee}(t'_7) &= \{\langle b \rangle_1\} \\ \Phi_{\vee}(t'_8) &= \{\langle b \rangle_1, \langle c \rangle_1\} \end{aligned}$$

If we focus on t'_7 because $\Delta_{7,a}(t'_7) > \Delta_{8,a}(t'_7)$ and its Φ_{\vee} -set is minimal, then $t'_7 \not\models \langle c \rangle_1 \models t'_8$ with $\langle c \rangle_1 \in \Phi_{\vee}(t'_8) \setminus \Phi_{\vee}(t'_7)$. As a consequence, $t_7 \not\models \langle a \rangle_1 \langle c \rangle_1 \models t_8$ where the value 1 decorating the a -diamond stems from $1 - \Delta_{8,a}(t'_7)$. ■

EXAMPLE 4.9. For the nodes t_1 and t_2 in Fig. 1(a), we have:

$$\begin{aligned} \Phi_{\vee}(t_1) &= \{\langle a \rangle_1, \langle a \rangle_{0.5} \langle b \rangle_1, \langle a \rangle_{0.5} \langle c \rangle_1\} \\ \Phi_{\vee}(t_2) &= \{\langle a \rangle_1, \langle a \rangle_{0.5} \langle b \rangle_1, \langle a \rangle_{0.5} \langle c \rangle_1, \langle a \rangle_1 (\langle b \rangle_1 \vee \langle c \rangle_1)\} \end{aligned}$$

The formulas with two diamonds and no disjunction are identical in the two sets, so their disjunction $\langle a \rangle_{0.5} \langle b \rangle_1 \vee \langle a \rangle_{0.5} \langle c \rangle_1$ is useless for discriminating between t_1 and t_2 . Indeed, such a formula is part of neither $\Phi_{\vee}(t_1)$ nor $\Phi_{\vee}(t_2)$. In contrast, their disjunction in which decorations of identical diamonds are summed up, i.e., $\langle a \rangle_1 (\langle b \rangle_1 \vee \langle c \rangle_1)$, is fundamental. It belongs only to $\Phi_{\vee}(t_2)$ because in the case of t_1 the b -transition and the c -transition depart from the same node, hence no probabilities can be added.

The fact that $\langle a \rangle_1 (\langle b \rangle_1 \vee \langle c \rangle_1) \in \Phi_{\vee}(t_2)$ is a distinguishing formula can be retrieved as follows. Starting from the two identically labeled transitions $t_1 \xrightarrow{a} \Delta_{1,a}$ and $t_2 \xrightarrow{a} \Delta_{2,a}$ where $\Delta_{1,a}(t'_1) = \Delta_{1,a}(t'_2) = 0.5 = \Delta_{2,a}(t'_2) = \Delta_{2,a}(t'_1)$ and $\Delta_{1,a}(t'_2) = \Delta_{1,a}(t'_1) = 0 = \Delta_{2,a}(t'_1) = \Delta_{2,a}(t'_2)$, we have:

$$\begin{aligned} \Phi_{\vee}(t'_1) &= \{\langle b \rangle_1, \langle c \rangle_1\} & \Phi_{\vee}(t'_2) &= \emptyset \\ \Phi_{\vee}(t'_2) &= \{\langle b \rangle_1\} & \Phi_{\vee}(t'_1) &= \{\langle c \rangle_1\} \end{aligned}$$

If we focus on t'_1 because $\Delta_{1,a}(t'_1) > \Delta_{2,a}(t'_1)$ and its Φ_{\vee} -set is minimal, then $t'_1 \not\models \langle b \rangle_1 \models t'_2$ with $\langle b \rangle_1 \in \Phi_{\vee}(t'_2) \setminus \Phi_{\vee}(t'_1)$ as well as $t'_1 \not\models \langle c \rangle_1 \models t'_2$ with $\langle c \rangle_1 \in \Phi_{\vee}(t'_2) \setminus \Phi_{\vee}(t'_1)$. As a consequence, $t_1 \not\models \langle a \rangle_1 (\langle b \rangle_1 \vee \langle c \rangle_1) \models t_2$ where the value 1 decorating the a -diamond stems from $1 - \Delta_{2,a}(t'_1)$. ■

EXAMPLE 4.10. For the nodes t_5 and t_6 in Fig. 1(c), we have:

$$\begin{aligned} \Phi_{\vee}(t_5) &= \{\langle a \rangle_1, \langle a \rangle_{0.25} \langle b \rangle_1, \langle a \rangle_{0.25} \langle c \rangle_1, \langle a \rangle_{0.5} (\langle b \rangle_1 \vee \langle c \rangle_1)\} \\ \Phi_{\vee}(t_6) &= \{\langle a \rangle_1, \langle a \rangle_{0.5} \langle b \rangle_1, \langle a \rangle_{0.5} \langle c \rangle_1\} \end{aligned}$$

The formulas with two diamonds and no disjunction are different in the two sets, so they are enough for discriminating between t_5 and t_6 . In contrast, the only formula with disjunction, which belongs to $\Phi_{\vee}(t_5)$, is useless because the probability decorating its a -diamond is equal to the probability decorating the a -diamond of each of the two formulas with two diamonds in $\Phi_{\vee}(t_6)$.

The fact that $\langle a \rangle_{0.5} \langle b \rangle_1 \in \Phi_{\vee}(t_6)$ is a distinguishing formula can be retrieved as follows. Starting from the two identically labeled transitions $t_5 \xrightarrow{a} \Delta_{5,a}$ and $t_6 \xrightarrow{a} \Delta_{6,a}$ where $\Delta_{5,a}(t'_5) = \Delta_{5,a}(t'_6) = 0.25$, $\Delta_{5,a}(t'') = 0.5 = \Delta_{6,a}(t'_6) = \Delta_{6,a}(t'')$, and $\Delta_{5,a}(t'_6) = 0 = \Delta_{6,a}(t'_5) = \Delta_{6,a}(t'_5)$, we have:

$$\begin{aligned} \Phi_{\vee}(t'_5) &= \{\langle b \rangle_1\} & \Phi_{\vee}(t'') &= \{\langle c \rangle_1\} \\ \Phi_{\vee}(t'_6) &= \{\langle b \rangle_1, \langle c \rangle_1\} & \Phi_{\vee}(t'') &= \emptyset \end{aligned}$$

Notice that t'' might be useless for discriminating purposes because it has the same probability in both distributions, so we exclude it. If we focus on t'_5 because $\Delta_{5,a}(t'_5) > \Delta_{6,a}(t'_5)$ and its Φ_{\vee} -set is minimal after the exclusion of t'' , then $t'_5 \not\models \langle b \rangle_1 \models t'_6$ with $\langle b \rangle_1 \in \Phi_{\vee}(t'_6) \setminus \Phi_{\vee}(t'_5)$, while no distinguishing formula is considered with respect to t'' as element of $\text{supp}(\Delta_{6,a})$ due to the exclusion of t'' itself. As a consequence, $t_5 \not\models \langle a \rangle_{0.5} \langle b \rangle_1 \models t_6$ where the value 0.5 decorating the a -diamond stems from $1 - (\Delta_{6,a}(t'_5) + p)$ with $p = \Delta_{6,a}(t'')$. The reason for subtracting the probability that t_6 reaches t'' after performing a is that $t'' \not\models \langle b \rangle_1$.

We conclude by observing that focussing on t'' as derivative with the minimum Φ_{\vee} -set is indeed problematic, because it would result in $\langle a \rangle_{0.5} \langle b \rangle_1$ when considering t'' as derivative of t_5 , but it would result in $\langle a \rangle_{0.5} (\langle b \rangle_1 \vee \langle c \rangle_1)$ when considering t'' as derivative of t_6 , with the latter formula not distinguishing between t_5 and t_6 . Moreover, when focussing on t'_5 , no formula ϕ' could have been found such that $t'_5 \not\models \phi' \models t''$ as $\Phi_{\vee}(t'') \subseteq \Phi_{\vee}(t'_5)$. ■

The last example shows that, in the general format for the PML_{\vee} distinguishing formula mentioned at the beginning of this subsection, i.e., $\langle a \rangle_{1 - (\Delta_{2,a}(t') + p)} \bigvee_{j \in J} \phi'_{2,j}$, the set J only contains any derivative of the second node different from t' to which the two distributions assign two *different* probabilities. No derivative of the two original nodes having the same probability in both distributions is taken into account even if its Φ_{\vee} -set is minimal – because it might be useless for discriminating purposes – nor is it included in J – because there might be no formula satisfied by this node when viewed as a derivative of the second node, which is not satisfied by t' . Furthermore, the value p is the probability that the second node reaches the excluded derivatives that do *not* satisfy $\bigvee_{j \in J} \phi'_{2,j}$; note that the first node reaches those derivatives with the same probability p .

We present two additional examples illustrating some technicalities of Def. 4.7. The former example shows the usefulness of the operator \bigvee and of the function *hplb* for selecting the right t' on the basis of the minimality of its Φ_{\vee} -set among the derivatives of the first node to which the first distribution assigns a probability greater than the second one. The latter example emphasizes the role played, for the same purpose as before, by formulas occurring in a Φ_{\vee} -set whose number of nested diamonds is not maximal.

EXAMPLE 4.11. For the nodes t_9 and t_{10} in Fig. 1(e), we have:

$$\begin{aligned}\Phi_V(t_9) &= \{\langle a \rangle_1, \langle a \rangle_{0.5} \langle b \rangle_1, \langle a \rangle_{0.5} \langle c \rangle_1\} \\ \Phi_V(t_{10}) &= \{\langle a \rangle_1, \langle a \rangle_{0.5} \langle b \rangle_1, \langle a \rangle_{0.5} \langle c \rangle_1, \langle a \rangle_{0.6} (\langle b \rangle_1 \vee \langle c \rangle_1)\}\end{aligned}$$

Starting from the two identically labeled transitions $t_9 \xrightarrow{a} \Delta_{9,a}$ and $t_{10} \xrightarrow{a} \Delta_{10,a}$ where $\Delta_{9,a}(t') = \Delta_{9,a}(t'') = 0.5$, $\Delta_{10,a}(t') = \Delta_{10,a}(t'') = 0.4$, $\Delta_{10,a}(t_{10}') = \Delta_{10,a}(t_{10}'') = 0.1$, and $\Delta_{9,a}(t_{10}') = \Delta_{9,a}(t_{10}'') = 0$, we have:

$$\begin{aligned}\Phi_V(t') &= \{\langle b \rangle_1, \langle c \rangle_1\} & \Phi_V(t'') &= \emptyset \\ \Phi_V(t_{10}') &= \{\langle b \rangle_1\} & \Phi_V(t_{10}'') &= \{\langle c \rangle_1\}\end{aligned}$$

If we focus on t'' because $\Delta_{9,a}(t'') > \Delta_{10,a}(t'')$ and its Φ_V -set is minimal, then $t'' \not\models \langle b \rangle_1 \equiv t'$ with $\langle b \rangle_1 \in \Phi_V(t') \setminus \Phi_V(t'')$, $t'' \not\models \langle b \rangle_1 \equiv t_{10}'$ with $\langle b \rangle_1 \in \Phi_V(t_{10}') \setminus \Phi_V(t'')$, and $t'' \not\models \langle c \rangle_1 \equiv t_{10}''$ with $\langle c \rangle_1 \in \Phi_V(t_{10}'') \setminus \Phi_V(t'')$. As a consequence, $t_9 \not\models \langle a \rangle_{0.6} (\langle b \rangle_1 \vee \langle c \rangle_1) \equiv t_{10}$ where the formula belongs to $\Phi_V(t_{10})$ and the value 0.6 decorating the a -diamond stems from $1 - \Delta_{10,a}(t'')$.

If \vee were used in place of $\dot{\vee}$, then in $\Phi_V(t_{10})$ we would also have formulas like $\langle a \rangle_{0.5} (\langle b \rangle_1 \vee \langle c \rangle_1)$ and $\langle a \rangle_{0.5} (\langle c \rangle_1 \vee \langle b \rangle_1)$. These are useless in that logically equivalent to other formulas already in $\Phi_V(t_{10})$ in which disjunction does not occur and, most importantly, would apparently augment the size of $\Phi_V(t_{10})$, an inappropriate fact in the case that t_{10} were a derivative of some other node instead of being the root of a tree.

If $hplb$ were not used, then in $\Phi_V(t_{10})$ we would also have formulas like $\langle a \rangle_{0.1} \langle b \rangle_1$, $\langle a \rangle_{0.4} \langle b \rangle_1$, $\langle a \rangle_{0.1} \langle c \rangle_1$, and $\langle a \rangle_{0.4} \langle c \rangle_1$, in which the probabilistic lower bounds of the a -diamonds are not maximal with respect to the satisfiability of formulas of that form by t_{10} ; those with maximal probabilistic lower bounds associated with a -diamonds are $\langle a \rangle_{0.5} \langle b \rangle_1$ and $\langle a \rangle_{0.5} \langle c \rangle_1$, which already belong to $\Phi_V(t_{10})$. In the case that t_9 and t_{10} were derivatives of two nodes under comparison instead of being the roots of two trees, the presence of those additional formulas in $\Phi_V(t_{10})$ may lead to focus on t_{10} instead of t_9 – for reasons that will be clear in Ex. 4.13 – thereby producing no distinguishing formula. ■

EXAMPLE 4.12. For the nodes t_{11} , t_{12} , t_{13} in Fig. 1(f), we have:

$$\begin{aligned}\Phi_V(t_{11}) &= \{\langle a \rangle_1\} \\ \Phi_V(t_{12}) &= \{\langle a \rangle_1, \langle a \rangle_1 \langle b \rangle_1\} \\ \Phi_V(t_{13}) &= \{\langle a \rangle_1, \langle a \rangle_{0.7} \langle b \rangle_1\}\end{aligned}$$

Let us view them as derivatives of other nodes, rather than roots of trees. The presence of formula $\langle a \rangle_1$ in $\Phi_V(t_{12})$ and $\Phi_V(t_{13})$ – although it has not the maximum number of nested diamonds in those two sets – ensures the minimality of $\Phi_V(t_{11})$ and hence that t_{11} is selected for building a distinguishing formula. If $\langle a \rangle_1$ were not in $\Phi_V(t_{12})$ and $\Phi_V(t_{13})$, then t_{12} and t_{13} could be selected, but no distinguishing formula satisfied by t_{11} could be obtained. ■

The criterion for selecting the right t' based on the minimality of its Φ_V -set has to take into account a further aspect related to *formulas without disjunctions*. If two derivatives – with different probabilities in the two distributions – have the same formulas without disjunctions in their Φ_V -sets, then a distinguishing formula for the two nodes will have disjunctions in it (see Exs. 4.9 and 4.11). In contrast, if the formulas without disjunctions are different between the two Φ_V -sets, then one of those formulas will tell the two derivatives apart (see Ex. 4.8).

A particular instance of the second case is the one in which for each formula without disjunctions in one of the two Φ_V -sets there is a variant in the other Φ_V -set – i.e., a formula without disjunctions that has the same format but may differ for the values of some probabilistic lower bounds – and vice versa. In this event, *regardless of the minimality* of the Φ_V -sets, it has to be selected the derivative such that (i) for each formula without disjunctions

in its Φ_V -set there exists a variant in the Φ_V -set of the other derivative such that the probabilistic lower bounds in the former formula are \leq than the corresponding bounds in the latter formula and (ii) at least one probabilistic lower bound in a formula without disjunctions in the Φ_V -set of the selected derivative is $<$ than the corresponding bound in the corresponding variant in the Φ_V -set of the other derivative. We say that the Φ_V -set of the selected derivative is a $(\leq, <)$ -variant of the Φ_V -set of the other derivative.

EXAMPLE 4.13. Let us view the nodes t_5 and t_6 in Fig. 1(c) as derivatives of other nodes, rather than roots of trees. Based on their Φ_V -sets shown in Ex. 4.10, we should focus on t_6 because $\Phi_V(t_6)$ contains fewer formulas. However, by so doing, we would be unable to find a distinguishing formula in $\Phi_V(t_5)$ that is not satisfied by t_6 . Indeed, if we look carefully at the formulas without disjunctions in $\Phi_V(t_5)$ and $\Phi_V(t_6)$, we note that they differ only for their probabilistic lower bounds: $\langle a \rangle_1 \in \Phi_V(t_6)$ is a variant of $\langle a \rangle_1 \in \Phi_V(t_5)$, $\langle a \rangle_{0.5} \langle b \rangle_1 \in \Phi_V(t_6)$ is a variant of $\langle a \rangle_{0.25} \langle b \rangle_1 \in \Phi_V(t_5)$, and $\langle a \rangle_{0.5} \langle c \rangle_1 \in \Phi_V(t_6)$ is a variant of $\langle a \rangle_{0.25} \langle c \rangle_1 \in \Phi_V(t_5)$. Therefore, we must focus on t_5 because $\Phi_V(t_5)$ contains formulas without disjunctions such as $\langle a \rangle_{0.25} \langle b \rangle_1$ and $\langle a \rangle_{0.25} \langle c \rangle_1$ having smaller bounds: $\Phi_V(t_5)$ is a $(\leq, <)$ -variant of $\Phi_V(t_6)$.

Consider now the nodes t_9 and t_{10} in Fig. 1(e), whose Φ_V -sets are shown in Ex. 4.11. If function $hplb$ were not used and hence $\Phi_V(t_{10})$ also contained $\langle a \rangle_{0.1} \langle b \rangle_1$, $\langle a \rangle_{0.4} \langle b \rangle_1$, $\langle a \rangle_{0.1} \langle c \rangle_1$, and $\langle a \rangle_{0.4} \langle c \rangle_1$, then the formulas without disjunctions in $\Phi_V(t_9)$ would no longer be equal to those in $\Phi_V(t_{10})$. More precisely, the formulas without disjunctions would be similar between the two sets, with those in $\Phi_V(t_{10})$ having smaller probabilistic lower bounds, so that we would erroneously focus on t_{10} . ■

Summing up, in the construction of the PML_V distinguishing formula mentioned at the beginning of this subsection, i.e., $\langle a \rangle_{1-(\Delta_{2,a}(t') + p)} \bigvee_{j \in J} \phi'_{2,j}$, the steps for choosing the derivative t' , on the basis of which each subformula $\phi'_{2,j}$ is then generated so that it is not satisfied by t' , are the following:

1. Consider only derivatives to which $\Delta_{1,a}$ assigns a probability greater than the one assigned by $\Delta_{2,a}$.
2. Within the previous set, eliminate all the derivatives whose Φ_V -sets have $(\leq, <)$ -variants.
3. Among the remaining derivatives, focus on one of those having a minimal Φ_V -set.

THEOREM 4.14. Let (T, A, \longrightarrow) be in RPT_f and $t_1, t_2 \in T$. Then $t_1 = t_2$ iff $t_1 \models \phi \iff t_2 \models \phi$ for all $\phi \in PML_V$. Moreover, if $t_1 \neq t_2$, then there exists $\phi \in PML_V$ distinguishing t_1 from t_2 such that $depth(\phi) \leq \max(height(t_1), height(t_2))$.

PROOF Given $t_1, t_2 \in T$, we proceed as follows:

- If $t_1 = t_2$, then obviously $t_1 \models \phi \iff t_2 \models \phi$ for all $\phi \in PML_V$.
- Assuming that $t_1 \neq t_2$, we show that there exists $\phi \in \Phi_V(t_1) \cup \Phi_V(t_2)$, which ensures that $depth(\phi) \leq \max(height(t_1), height(t_2))$, such that it is not the case that $t_1 \models \phi \iff t_2 \models \phi$ by proceeding by induction on $height(t_1) \in \mathbb{N}$. The proof is similar to the one of Thm. 4.4, in particular in the cases $height(t_1) = 0$ and $height(t_1) = n + 1$ with $init(t_1) \neq init(t_2)$ it benefits from the presence of $\{\langle a_i \rangle_1 \mid i \in I\}$ in $\Phi_V(t)$ as of Def. 4.7. However, it changes as follows before the application of the induction hypothesis in the case $height(t_1) = n + 1$ with $init(t_1) = init(t_2) \neq \emptyset$ and $t_1 \xrightarrow{a} \Delta_{1,a}$, $t_2 \xrightarrow{a} \Delta_{2,a}$, and $\Delta_{1,a} \neq \Delta_{2,a}$ for some $a \in init(t_1)$. Let $supp_a = supp(\Delta_{1,a}) \cup supp(\Delta_{2,a})$, which can be partitioned into $supp_{a,\neq} = \{t' \in supp_a \mid \Delta_{1,a}(t') \neq \Delta_{2,a}(t')\}$

and $\text{supp}_{a,=} = \{t' \in \text{supp}_a \mid \Delta_{1,a}(t') = \Delta_{2,a}(t')\}$ with $|\text{supp}_{a,=}| \geq 2$ because $\Delta_{1,a} \neq \Delta_{2,a}$ and $|\text{supp}_{a,=}| \geq 0$. We recall that $\Phi_V(t'')$ is a $(\leq, <)$ -variant of $\Phi_V(t')$ iff:

- for each formula without disjunctions in one of the two Φ_V -sets there is a variant in the other Φ_V -set – i.e., a formula without disjunctions that has the same format but may differ for the values of some probabilistic lower bounds – and vice versa (this means that there exists a bijection between the formulas without disjunctions in the two Φ_V -sets, because the maximality of the probabilistic lower bounds in a Φ_V -set implies the existence of at most one formula with a given format in the Φ_V -set);
- for each formula without disjunctions in $\Phi_V(t'')$ there exists a variant in $\Phi_V(t')$ such that the probabilistic lower bounds in the former formula are \leq than the corresponding bounds in the latter formula;
- at least one probabilistic lower bound in a formula without disjunctions in $\Phi_V(t'')$ is $<$ than the corresponding bound in the corresponding variant in $\Phi_V(t')$.

Among all the nodes in $\text{supp}_{a,\neq}$, there exists one denoted by t' such that, for all $t'' \in \text{supp}_{a,\neq} \setminus \{t'\}$, $\Phi_V(t'')$ is not a $(\leq, <)$ -variant of $\Phi_V(t')$ as we now prove by proceeding by induction on $|\text{supp}_{a,\neq}| \in \mathbb{N}_{\geq 2}$:

- If $|\text{supp}_{a,\neq}| = 2$ – hence $\text{supp}_{a,\neq} = \{t', t''\}$ – then trivially at least one of $\Phi_V(t')$ and $\Phi_V(t'')$ is not a $(\leq, <)$ -variant of the other.
- Let $|\text{supp}_{a,\neq}| = n+1$ for some $n \in \mathbb{N}_{\geq 2}$ and suppose that the result holds for each subset of $\text{supp}_{a,\neq}$ of cardinality between 2 and n . Assuming that $\text{supp}_{a,\neq} = \{t'_1, t'_2, \dots, t'_{n+1}\}$, we denote by t' the node in $\text{supp}_{a,\neq} \setminus \{t'_{n+1}\}$ that, by the induction hypothesis, enjoys the property over that subset. There are two cases:
 - If $\Phi_V(t'_{n+1})$ is not a $(\leq, <)$ -variant of $\Phi_V(t')$ either, then t' enjoys the property over the entire set $\text{supp}_{a,\neq}$.
 - Suppose that $\Phi_V(t'_{n+1})$ is a $(\leq, <)$ -variant of $\Phi_V(t')$, which implies that $\Phi_V(t')$ cannot be a $(\leq, <)$ -variant of $\Phi_V(t'_{n+1})$. From the fact that, for all $t'' \in \text{supp}_{a,\neq} \setminus \{t', t'_{n+1}\}$, $\Phi_V(t'')$ is not a $(\leq, <)$ -variant of $\Phi_V(t')$ by the induction hypothesis, it follows that $\Phi_V(t'')$ is not a $(\leq, <)$ -variant of $\Phi_V(t'_{n+1})$. Indeed, for each such t'' the set $\Phi_V(t'')$ contains at least a formula without disjunctions that is not a variant of any formula without disjunctions in $\Phi_V(t')$, or all formulas without disjunctions in $\Phi_V(t'')$ are identical to formulas without disjunctions in $\Phi_V(t')$, hence this holds true with respect to $\Phi_V(t'_{n+1})$ too, given that $\Phi_V(t'_{n+1})$ is a $(\leq, <)$ -variant of $\Phi_V(t')$. As a consequence, t'_{n+1} enjoys the property over the entire set $\text{supp}_{a,\neq}$.

Within the set of all the nodes in $\text{supp}_{a,\neq}$ enjoying the property above, we select one with a minimal Φ_V -set, which we denote by t'_{\min} . Suppose that $\Delta_{1,a}(t'_{\min}) > \Delta_{2,a}(t'_{\min})$ and let $t'_{2,j}$ be an arbitrary node belonging to $\text{supp}_{a,\neq,2} = (\text{supp}_{a,\neq} \setminus \{t'_{\min}\}) \cap \text{supp}(\Delta_{2,a})$. By the induction hypothesis, from $t'_{2,j} \neq t'$ it follows that there exists $\phi'_{2,j} \in \Phi_V(t'_{\min}) \cup \Phi_V(t'_{2,j})$ such that it is not the case that $t'_{\min} \models \phi'_{2,j} \iff t'_{2,j} \models \phi'_{2,j}$. In particular, it holds that $t'_{\min} \not\models \phi'_{2,j} \implies t'_{2,j} \models \phi'_{2,j}$ because $\phi'_{2,j} \in \Phi_V(t'_{2,j})$, as can be seen by considering the following two cases based on the fact that $\Phi_V(t'_{2,j})$ is not a $(\leq, <)$ -variant of $\Phi_V(t'_{\min})$:

- If at least one formula without disjunctions in $\Phi_V(t'_{2,j})$ is not a variant of any formula without disjunctions in $\Phi_V(t'_{\min})$, then such a formula can be taken as $\phi'_{2,j}$ given the maximality of the probabilistic lower bounds of any basic formula in $\Phi_V(t'_{\min})$.
- If all basic formulas in $\Phi_V(t'_{2,j})$ are identical to basic formulas in $\Phi_V(t'_{\min})$, then $\Phi_V(t'_{2,j})$ must contain some more formulas

(with disjunctions) not in $\Phi_V(t'_{\min})$ given the minimality of the latter set, otherwise we would have selected $t'_{2,j}$ in place of t'_{\min} . One of the additional formulas (with disjunctions) in $\Phi_V(t'_{2,j})$ can be taken as $\phi'_{2,j}$.

Letting $\text{supp}_{a,=,\neq} = \{t' \in \text{supp}_{a,=} \mid t' \not\models \bigvee_{t'_{2,j} \in \text{supp}_{a,\neq,2}} \phi'_{2,j}\}$ as well as $p_{\neq} = \Delta_{2,a}(\text{supp}_{a,=,\neq}) = \Delta_{1,a}(\text{supp}_{a,=,\neq})$, we have that $t_1 \not\models \langle a \rangle_{1-(\Delta_{2,a}(t'_{\min})+p_{\neq})} \bigvee_{t'_{2,j} \in \text{supp}_{a,\neq,2}} \phi'_{2,j} \models t_2$ because $1 - (\Delta_{2,a}(t'_{\min}) + p_{\neq}) > 1 - (\Delta_{1,a}(t'_{\min}) + p_{\neq})$ and the maximum probabilistic lower bound for which t_1 satisfies a formula of that form cannot exceed $1 - (\Delta_{1,a}(t'_{\min}) + p_{\neq})$. The PML_V distinguishing formula above may not be in $\Phi_V(t_2)$, but it is logically implied by, or equivalent to, a distinguishing formula in $\Phi_V(t_2)$ for the following reasons:

- Each $t'_{2,j}$ belongs to $\text{supp}(\Delta_{2,a})$.
- Each $\phi'_{2,j}$ belongs to $\Phi_V(t'_{2,j})$.
- The probabilistic lower bound $1 - (\Delta_{2,a}(t'_{\min}) + p_{\neq})$ is equal to $\sum_{t'_{2,j} \in \text{supp}_{a,\neq,2}} \Delta_{2,a}(t'_{2,j}) + \Delta_{2,a}(\text{supp}_{a,=,\neq})$, so in the PML_V distinguishing formula it is sufficient to replace $\bigvee_{t'_{2,j} \in \text{supp}_{a,\neq,2}} \phi'_{2,j}$ with $\bigvee_{t' \in \text{supp}_{a,\neq,2} \cup \text{supp}_{a,=,\neq}} \phi_{t'}$ where:
 - $\phi_{t'} = \phi'_{2,j}$ if $t' = t'_{2,j}$ for some j ;
 - $\phi_{t'} = \phi' \in \Phi_V(t')$ if $t' \neq t'_{2,j}$ for all j , where $\phi' \implies \phi'_{2,j}$ for some j and the existence of such a ϕ' in $\Phi_V(t')$ stems from $t' \in \text{supp}_{a,=,\neq}$, i.e., $t' \models \bigvee_{t'_{2,j} \in \text{supp}_{a,\neq,2}} \phi'_{2,j}$. ■

4.4 PML_Δ Characterizes \sim_{PB} : A Direct Proof for Discrete Systems

By adapting the proof of Thm. 4.14 consistently with the proof of Thm. 4.3, in our setting we can also demonstrate that PML_Δ characterizes \sim_{PB} by working directly on *discrete* state spaces. The idea is to obtain $t_1 \models \langle a \rangle_{\Delta_{1,a}(t') + p} \bigwedge_{j \in J} \phi'_{2,j} \not\models t_2$.

To this purpose, for any node t of finite height we define the set $\Phi_\Delta(t)$ of PML_Δ formulas satisfied by t featuring, in addition to maximal probabilistic lower bounds and diamonds arising only from transitions of t as for $\Phi_V(t)$, conjunctions that (i) stem only from transitions departing from the *same node* in the support of a distribution reached by t and (ii) are preceded by a diamond decorated with the *sum* of the probabilities assigned by that distribution to that node and other nodes with the *same transitions* considered for that node. Formally, given t having transitions of the form $t \xrightarrow{a_i} \Delta_i$ with $\text{supp}(\Delta_i) = \{t'_{i,j} \mid j \in J_i\}$ and $i \in I \neq \emptyset$, we let:

$$\begin{aligned} \Phi_\Delta(t) = & \{ \langle a_i \rangle_1 \mid i \in I \} \\ & \cup \bigcup_{i \in I} \text{splb}(\{ \langle a_i \rangle_{\Delta_i(t'_{i,j})} \bigwedge_{k \in K'} \phi'_{i,j,k} \mid \emptyset \neq K' \subseteq K_{i,j}, \\ & \quad t'_{i,j} \in \text{supp}(\Delta_i), \phi'_{i,j,k} \in \Phi_\Delta(t'_{i,j}) \}) \end{aligned}$$

where $\{ \}$ and $\{ \}$ are multiset parentheses, $K_{i,j}$ is the index set for $\Phi_\Delta(t'_{i,j})$, and function splb merges all formulas possibly differing only for the probabilistic lower bound decorating their initial a_i -diamond by summing up those bounds (notice that such formulas stem from different nodes in $\text{supp}(\Delta_i)$).

We now provide some examples illustrating the technicalities of the definition above, as well as the fact that a good criterion for choosing t' occurring in the PML_Δ distinguishing formula at the beginning of this subsection is the *maximality* of the Φ_Δ -set.

EXAMPLE 4.15. In Fig. 1(b), the multiset giving rise to $\Phi_\Delta(t_3)$ contains two occurrences of $\langle a \rangle_{0.2} \langle b \rangle_1$ and two occurrences of $\langle a \rangle_{0.1} \langle b \rangle_1$, which are merged into $\langle a \rangle_{0.6} \langle b \rangle_1$ by function splb . Likewise, the multiset behind $\Phi_\Delta(t_4)$ contains formulas $\langle a \rangle_{0.1} \langle b \rangle_1$, $\langle a \rangle_{0.3} \langle b \rangle_1$, and $\langle a \rangle_{0.2} \langle b \rangle_1$, which are merged into $\langle a \rangle_{0.6} \langle b \rangle_1$. ■

EXAMPLE 4.16. For the nodes t_1 and t_2 in Fig. 1(a), we have:

$$\begin{aligned}\Phi_{\wedge}(t_1) &= \{\langle a \rangle_1, \langle a \rangle_{0.5} \langle b \rangle_1, \langle a \rangle_{0.5} \langle c \rangle_1, \langle a \rangle_{0.5} (\langle b \rangle_1 \wedge \langle c \rangle_1)\} \\ \Phi_{\wedge}(t_2) &= \{\langle a \rangle_1, \langle a \rangle_{0.5} \langle b \rangle_1, \langle a \rangle_{0.5} \langle c \rangle_1\}\end{aligned}$$

The conjunction $\langle a \rangle_{0.5} \langle b \rangle_1 \wedge \langle a \rangle_{0.5} \langle c \rangle_1$ is useless for discriminating between t_1 and t_2 – it is part of neither $\Phi_{\wedge}(t_1)$ nor $\Phi_{\wedge}(t_2)$ – while $\langle a \rangle_{0.5} (\langle b \rangle_1 \wedge \langle c \rangle_1)$ is the only distinguishing formula and belongs only to $\Phi_{\wedge}(t_1)$, because in the case of t_2 the b -transition and the c -transition depart from two different nodes. Starting from the two identically labeled transitions $t_1 \xrightarrow{a} \Delta_{1,a}$ and $t_2 \xrightarrow{a} \Delta_{2,a}$ where $\Delta_{1,a}(t'_1) = \Delta_{1,a}(t''_1) = 0.5 = \Delta_{2,a}(t'_2) = \Delta_{2,a}(t''_2)$ and $\Delta_{1,a}(t'_2) = \Delta_{1,a}(t''_2) = 0 = \Delta_{2,a}(t'_1) = \Delta_{2,a}(t''_1)$, we have:

$$\begin{aligned}\Phi_{\wedge}(t'_1) &= \{\langle b \rangle_1, \langle c \rangle_1\} & \Phi_{\wedge}(t'_1) &= \emptyset \\ \Phi_{\wedge}(t'_2) &= \{\langle b \rangle_1\} & \Phi_{\wedge}(t'_2) &= \{\langle c \rangle_1\}\end{aligned}$$

If we focus on t'_1 because $\Delta_{1,a}(t'_1) > \Delta_{2,a}(t'_1)$ and its Φ_{\wedge} -set is maximal, then $t'_1 \models \langle c \rangle_1 \not\models t'_2$ with $\langle c \rangle_1 \in \Phi_{\wedge}(t'_1) \setminus \Phi_{\wedge}(t'_2)$ as well as $t'_1 \models \langle b \rangle_1 \not\models t'_2$ with $\langle b \rangle_1 \in \Phi_{\wedge}(t'_1) \setminus \Phi_{\wedge}(t'_2)$. As a consequence, $t_1 \models \langle a \rangle_{0.5} (\langle b \rangle_1 \wedge \langle c \rangle_1) \not\models t_2$ where the value 0.5 decorating the a -diamond stems from $\Delta_{1,a}(t'_1)$. ■

As far as the other two variables occurring in the PML_{\wedge} distinguishing formula at the beginning of this subsection are concerned, J only contains any derivative of the second node different from t' to which the two distributions assign two *different* probabilities, while p is the probability of reaching derivatives having the *same* probability in both distributions that *satisfy* $\bigwedge_{j \in J} \phi'_{2,j}$. Moreover, when selecting t' , we have to leave out all the derivatives whose Φ_{\wedge} -sets have $(\leq, <)$ -variants.

THEOREM 4.17. Let (T, A, \longrightarrow) be in RPT_f and $t_1, t_2 \in T$. Then $t_1 = t_2$ iff $t_1 \models \phi \iff t_2 \models \phi$ for all $\phi \in \text{PML}_{\wedge}$. Moreover, if $t_1 \neq t_2$, then there exists $\phi \in \text{PML}_{\wedge}$ distinguishing t_1 from t_2 such that $\text{depth}(\phi) \leq \max(\text{height}(t_1), \text{height}(t_2))$.

PROOF Similar to that of Thm. 4.14, with these differences:

- We select t'_{\max} as one of the nodes with maximal Φ_{\wedge} -set in $\text{supp}_{a, \neq}$ having no $(\leq, <)$ -variants.
- It holds that $t'_{\max} \models \phi'_{2,j} \not\models t'_{2,j}$ for all $t'_{2,j} \in \text{supp}_{a, \neq, 2}$ because $\phi'_{2,j} \in \Phi_{\wedge}(t'_{\max})$ thanks to the maximality of $\Phi_{\wedge}(t'_{\max})$.
- Letting $\text{supp}_{a, =, \models} = \{t' \in \text{supp}_{a, =} \mid t' \models \bigwedge_{t'_{2,j} \in \text{supp}_{a, \neq, 2}} \phi'_{2,j}\}$ as well as $p_{\models} = \Delta_{1,a}(\text{supp}_{a, =, \models}) = \Delta_{2,a}(\text{supp}_{a, =, \models})$, we have that $t_1 \models \langle a \rangle_{\Delta_{1,a}(t'_{\max}) + p_{\models}} \bigwedge_{t'_{2,j} \in \text{supp}_{a, \neq, 2}} \phi'_{2,j} \not\models t_2$ because $\Delta_{1,a}(t'_{\max}) + p_{\models} > \Delta_{2,a}(t'_{\max}) + p_{\models}$ and the maximum probabilistic lower bound for which t_2 satisfies a formula of that form cannot exceed $\Delta_{2,a}(t'_{\max}) + p_{\models}$.
- The PML_{\wedge} distinguishing formula is in $\Phi_{\wedge}(t_1)$ due to *splb*. ■

5. Conclusions

In this paper, we have studied modal logic characterizations of bisimilarity over reactive probabilistic processes. Starting from previous work by Larsen and Skou [12] (who provided a characterization based on an extension of Hennessy-Milner modal logic where diamonds are decorated with probabilistic lower bounds) and by Desharnais, Edalat, and Panangaden [5] (who showed that negation is not necessary), we have proved that conjunction can be replaced by disjunction without having to reintroduce negation. Thus, in the reactive probabilistic setting, conjunction and disjunction are *interchangeable* to characterize (bi)simulation equivalence, while they are *both necessary* for simulation preorder [6]. As a side result, using the same proof technique we have provided alternative proofs of the expressiveness of the logics $\text{PML}_{\neg \wedge}$ and PML_{\wedge} .

The intuition behind our result for PML_{\vee} is that from a conjunctive distinguishing formula it is often possible to derive a disjunctive one by suitably increasing some probabilistic lower bounds. This corresponds to summing up the probabilities of reaching certain states that are in the support of a target distribution. In fact, a state of an RPLTS can be given a semantics as a reactive probabilistic tree, and hence it is characterized by the countable set of formulas (approximated by the Φ_{\vee} -set) obtained by doing finite visits of the tree.

On the application side, the PML_{\vee} -based characterization of bisimilarity over reactive probabilistic processes may help to prove a conjecture in [3] about the discriminating power of three different testing equivalences respectively using reactive probabilistic tests, fully nondeterministic tests, and nondeterministic and probabilistic tests. Many examples lead to conjecture that testing equivalence based on nondeterministic and probabilistic tests has the same discriminating power as bisimilarity. Given two \sim_{PB} -inequivalent reactive probabilistic processes, the idea of the tentative proof is to build a distinguishing nondeterministic and probabilistic test from a distinguishing PML_{\wedge} formula. One of the main difficulties with carrying out such a proof, i.e., the fact that choices within tests fit well together with disjunction rather than conjunction, may be overcome by starting from a distinguishing PML_{\vee} formula.

References

- [1] P. Aczel and N. Mendler. A final coalgebra theorem. In *Proc. CTCS 1989*, volume 389 of *LNCS*, pages 357–365. Springer, 1989.
- [2] C. Baier and M. Kwiatkowska. Domain equations for probabilistic processes. *Mathematical Structures in Computer Science*, 10:665–717, 2000.
- [3] M. Bernardo, D. Sangiorgi, and V. Vignudelli. On the discriminating power of testing equivalences for reactive probabilistic systems: Results and open problems. In *Proc. QEST 2014*, volume 8657 of *LNCS*, pages 281–296. Springer, 2014.
- [4] Y. Deng and H. Wu. Modal characterisations of probabilistic and fuzzy bisimulations. In *Proc. ICFEM 2014*, volume 8829 of *LNCS*, pages 123–138. Springer, 2014.
- [5] J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labelled Markov processes. *Information and Computation*, 179:163–193, 2002.
- [6] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating labelled Markov processes. *Information and Computation*, 184:160–200, 2003.
- [7] R.J. van Glabbeek. The linear time – branching time spectrum I. In *Handbook of Process Algebra*, pages 3–99. Elsevier, 2001.
- [8] R.J. van Glabbeek, S.A. Smolka, and B. Steffen. Reactive, generative and stratified models of probabilistic processes. *Information and Computation*, 121:59–80, 1995.
- [9] M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM*, 32:137–162, 1985.
- [10] B. Jacobs and A. Sokolova. Exemplaric expressivity of modal logics. *Journal of Logic and Computation*, 20:1041–1068, 2010.
- [11] R.M. Keller. Formal verification of parallel programs. *Communications of the ACM*, 19:371–384, 1976.
- [12] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.
- [13] P. Panangaden. Probabilistic bisimulation. In *Advanced Topics in Bisimulation and Coinduction*, pages 300–334. CUP, 2011.
- [14] J.J.M.M. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.
- [15] E.P. de Vink and J.J.M.M. Rutten. Bisimulation for probabilistic transition systems: A coalgebraic approach. *Theoretical Computer Science*, 221:271–293, 1999.
- [16] J. Worrell. On the final sequence of a finitary set functor. *Theoretical Computer Science*, 338:184–199, 2005.